

A FORMAL APPROACH TO MODELING THE MEMORY OF A LIVING ORGANISM

DAN P. GURALNIK

ABSTRACT. We consider a living organism as an observer of the evolution of its environment recording sensory information about the state space X of the environment in real time. Sensory information is sampled and then processed on two levels. On the biological level, the organism serves as an *evaluation mechanism* of the subjective relevance of the incoming data to the observer: the observer assigns *excitation values* to events in X it could recognize using its sensory equipment. On the algorithmic level, sensory input is used for updating a database – the memory of the observer – whose purpose is to serve as a geometric/combinatorial model of X , whose nodes are weighted by the excitation values produced by the evaluation mechanism. These values serve as a guidance system for deciding how the database should transform as observation data mounts.

We define a searching problem for the proposed model and discuss the model's flexibility and its computational efficiency, as well as the possibility of implementing it as a dynamic network of neuron-like units. We show how various easily observable properties of the human memory and thought process can be explained within the framework of this model. These include: reasoning (with efficiency bounds), errors, temporary and permanent loss of information. We are also able to define general learning problems in terms of the new model, such as the language acquisition problem.

*Dedicated to the memory of my father,
Peter J. Guralnik, and all his mice and rats.*

1. INTRODUCTION

1.1. General considerations. The structure of memory in living organisms is generally perceived as extremely complex and extremely efficient at the same time. Complex, because the sheer number of physiological structure elements comprising the nervous system of a *rattus rattus*, say – not to mention *homo sapiens* – seems to exclude the possibility of direct piece-by-piece analysis. Efficient, because of the capability to respond to unpredicted input signals in real time, frequently with desirable outcome.

Different species seem to exhibit different capacities for this kind of thinking, and the reason for this remains unclear. By this we mean that, although biologists have a lot to say about the evolution of the central nervous system, and in spite of our ability to roughly map the brain and say which parts of it are in charge of which functions, one major problem remains completely open: *what is the algorithmic structure underlying the phenomenon we call 'intelligence'?*

Key words and phrases. Memory, database, learning, natural language acquisition, poc-set, poc-morphism, median graph, median morphism.

Answering this question adequately is one of the dream goals of the field of Artificial Intelligence, whose ultimate goal – one should naïvely guess – is to construct machines capable of humanoid reasoning, however much faster and more accurate.

Developing mathematical tools for describing the principles governing our thought process and the development of our minds may turn out useful for improving our understanding of learning and teaching, and may seriously impact psychology. Describing the relationship between language formation and acquisition, structure of memory, logical thinking and psychology may finally be within our grasp.

Considering the wide variety of applications it becomes desirable to formulate the answer to our main question in a way that is independent of its physical realization (e.g. the biophysics of the brain). We will refer to this as the ‘*invariance principle*’.

The invariance principle defines an objective measure of adequacy for our attempted answers. Indeed, if invariance is observed, then whenever an abstract model of the mind is offered, it becomes possible to study multiple realizations, which gives us the chance to search for the most efficient ones; if no efficient realization exists, we discard the model. On the other hand, if a model is tied to a specific physical realization, then its predictive power is automatically limited by our understanding of that realization. The latter usually being incomplete also puts us in jeopardy of using an intrinsically inconsistent model for predictions.

In this work we introduce an approach based on considering memory as an algorithmic structure: a database, where data is stored, together with a set of algorithms in charge of maintaining the structure and retrieving stored information. By ‘memory’ we mean the broadest possible interpretation of the term: all information observed and retained by the living organism in the course of its lifetime, together with the procedures handling all this data.

Though contributing to the feeling that analysis of such a ‘memory’ is hopelessly hard, this interpretation is a necessity dictated by the invariance principle. In the case of animals (including humans) it was Evolution that shaped our memory management system and distributed its functions among the many scales and subsystems of our organisms, solving this problem on the way, in the course of billions of years. It is then only reasonable to assume that when it comes to producing a *homunculus* of our own making, we will need to face the same challenges. Therefore there is no point in restricting our theoretical framework.

It is unclear how memory structures of living organisms are maintained and what principles are responsible for the seeming high efficiency of their recording and retrieval algorithms. In the preceding sentence, the word ‘seeming’ is used because, in contrast to the general feeling of awed amazement at the capabilities of human and animal brains, frequent everyday observations of humans strongly suggest that said brains possess some rather discouraging inherent flaws. For example, trouble recovering useful information while recalling seemingly irrelevant data with ease (“*what was that formula for $\cos 2\theta$?*” vs. “*can not get that idiotic tune out of my head!*”), failure to recognize phenomena that are supposedly well understood (“*why didn’t *I* think about that?!*”), as well as difficulties in processing logically complex statements (e.g. *he thinks that she thinks that*

he knows that she heard that he thinks she believes that he loves someone else), – are all everyday common examples of what we normally perceive as ‘glitches’ of our memory system. We feel that if only we could merge ourselves with a powerful computer, all these inaccuracies will be gone.

One possible reason for this feeling is that we tend to see our memory management structure as the result of a very successful evolutionary process. This idea makes it too easy for us to focus only on the desirable manifestations of memory structuring when trying to replicate nature’s achievements, while all the undesirable effects we observe on a daily basis are written off as malfunctions of a complicated analog system (our organism), occurring due to fatigue and/or some other physiological impairment.

On the other hand, one could also argue that the high frequency of malfunctions is the manifestation of a set of principles governing the way in which our memory resources are managed. Indeed, since *any* organizing principle in fact constitutes a restriction on the set of admissible structures (that is, some database structures will never be realized because they violate said principle), it is reasonable to expect many admissible structures to react inadequately to some unpredictable situations presented by the environment. Therefore, the optimist will regard all the failings of our memory storage system as hints to how such a system is structured, and the quality of a model should be judged not only by its computational efficiency, but more by its ability to explain the functional role of errors in the system, how errors occur, which errors occur more frequently than others, etc.

1.2. Stating the problem. It is time to state our goals in a more committing and formal fashion. We consider the organism’s memory as a record of its observations of the evolution of the environment (participation by the organism is not ruled out). The environment as a whole has an associated space of states – denoted henceforth by X , – and every organism \mathcal{O} maintains a database Γ whose structure and content correspond to that observer’s perception of properties of X , as determined by the available sensory equipment. Thus, we consider the organism \mathcal{O} as a mediator between the environment and the database Γ .

Recall that a database normally has the structure of a graph (or *network*), with nodes carrying additional information (or *content*). Updating the database may involve altering the content of a node or nodes, as well as adding new nodes or erasing redundant ones.

When an organism \mathcal{O} with memory structure Γ makes an observation about X (receives new input from the environment), this observation will be evaluated (this means Γ is being read), possibly resulting in an updating procedure and replacing Γ by a new structure Γ' . Mathematically speaking, this means our model should have an underlying structure of a category \mathcal{C} , whose objects belong to the class of admissible databases, and the morphisms (also known as *arrows*) describe the various possibilities of transforming one structure into the other. The structure (e.g. algebraic, topological, other...) of this category corresponds to the idea of a set of ‘governing principles’ we have just discussed.

Recall that a category \mathcal{C} consists of a class of objects $\text{ob}\mathcal{C}$, a set $\mathcal{C}(A, B)$ of morphisms $f : A \rightarrow B$, defined for every ordered pair of objects $A, B \in \text{ob}\mathcal{C}$, and an operation

(called composition), defined for each ordered triple A, B, C of objects –

$$\left\{ \begin{array}{ccc} \circ : \mathcal{C}(A, B) \times \mathcal{C}(B, C) & \longrightarrow & \mathcal{C}(A, C) \\ f \times g & \mapsto & g \circ f \end{array} \right.$$

Composable morphisms are said to be *compatible*. All the above must satisfy the following requirements:

- For every $A \in \text{ob}\mathcal{C}$, the set $\mathcal{C}(A, A)$ contains a distinguished element denoted id_A ;
- For every $A, B \in \text{ob}\mathcal{C}$ and any $f \in \mathcal{C}(A, B)$ one has $f \circ \text{id}_A = \text{id}_B \circ f = f$;
- For every triple of compatible morphisms f, g, h , one has $h \circ (g \circ f) = (h \circ g) \circ f$.

Contemplating the meaning of the axioms of a category in our context, one may regard a morphism between two objects (databases) as a means for comparing them. One may want to measure the improvement resulting from updating Γ into Γ' (see above), or one could be interested in measuring the differences in how two distinct observers perceive a common environment. Finally, one may want to measure the discrepancy between an observer's perception and the objective reality presented by the environment. Whether or not these are possible in any useful sense depends on the structure of the category underlying our modeling method.


The motivation for our construction has two sources: one is Shannon's idea of *entropy* introduced in [Sha48], and the other is the idea of *spaces with walls* introduced by Haglund and Paulin in [HP98]. The marriage of the two produces the information-theoretic approach to constructing databases based on binary observations, which we shall describe right now.

It is natural to assume the state space X of the observed environment is a topological space endowed with the corresponding Borel σ -algebra \mathcal{B} and a probability measure μ , making it possible to consider the probability of an event in X . Let us fix a moment in time and an observer \mathcal{O} and assume that the sensory equipment available to \mathcal{O} produces binary output. Assume there are only finitely many sensors available to \mathcal{O} . Imagine an angel (as opposed to the notion of a daemon, frequently used in the literature to explain various notions of entropy) who is in charge of recording the output from these sensors and keeping it in order.

The angel is absolutely objective: it does not have a preference to any kind of data generated by the sensors, and its sole and sacred responsibility is to record the data as accurately as possible. The angel is all-knowing: for each sensor, it knows precisely which inputs (states of the environment) produce which output for that sensor.

Thus, in the angel's notebook, each sensor will correspond to a pair of complementary subsets of X , while the totality of all information that \mathcal{O} is capable of producing then becomes a finite family H of subsets of X , which is closed under complementation. Even angels have limited powers, so $H \subseteq \mathcal{B}$.

Associated with H is a partition of X : given $x, y \in X$, write $x \stackrel{H}{\sim} y$ if and only if $x \in h$ implies $y \in h$ for every $h \in H$. The relation $(\stackrel{H}{\sim})$ is then an equivalence relation whose induced partition (denote it by $\mathcal{P}(H)$) is the join (the coarsest common refinement) of the binary partitions corresponding to the individual sensors. It is important that

FIGURE 1. Observing a compass (1). Two versions of the system H of example 1.1.

the angel keep this partition on the record: the observer \mathcal{O} is incapable of discerning between two events belonging to the same element of this partition. For this reason, the elements of $\mathcal{P}(H)$ will be called *H-visible states*.

From the point of view of the angel, it is now possible to determine the amount of information that \mathcal{O} has about X : this task is equivalent to computing Shannon's entropy of $\mathcal{P}(H)$. Unfortunately, this result is not meaningful for the observer \mathcal{O} : Shannon's entropy is approximated by the minimum expectation of the number of *arbitrary* binary questions one needs to ask in order to determine the position of a point of X with respect to $\mathcal{P}(H)$, but \mathcal{O} only has the questions from H at his disposal, so the minimum computed by \mathcal{O} may end up much higher than that computed by the all-knowing angel, who is surely capable of asking *any* question from the list \mathcal{B} .

Despite the differences in computational ability, both the observer and the angel are interested in having the probabilities $\mu(P)$ of each element $P \in \mathcal{P}(H)$ recorded somewhere. Thus, we have content for our database, but no graph to put it in. This is easily repaired. Define a graph Γ_H to have $\mathcal{P}(H)$ for its set of vertices, where two vertices $P, Q \in \mathcal{P}(H)$ are joined by an edge if and only if there exists precisely one $h \in H$ satisfying $P \subseteq h$ and $Q \subseteq h^c$ (here and on, $h^c = X \setminus h$). This construction is just a special case of the construction of the graph dual to a space with walls. Note that Γ_H is necessarily connected and bipartite. Perhaps it is time for an example:

Example 1.1 (observing a compass, part 1). Consider a person \mathcal{O} observing a compass. The space of states of the needle of the compass can be modeled by the unit circle $X = \mathbb{S}^1$, thought of as a subset of the complex numbers \mathbb{C} , with the number $\mathbf{1}$ corresponding to precise North, and \mathbf{i} corresponding to West.

Now, imagine our observer being able to ask the questions “Is the needle pointing North/South/West/East?” (and their complements, of course), and let us model the




FIGURE 2. Observing a compass (1). The graphs Γ_H for the two types of the system H from example 1.1 – compare with fig. 1.

positive answer sets – denoted by N, S, W and E , respectively – by open intervals:

$$\begin{aligned} N &= \left\{ e^{i\theta} \mid \theta \in (-\epsilon, \epsilon) \right\} \\ S &= \left\{ e^{i\theta} \mid \theta \in (\pi - \epsilon, \pi + \epsilon) \right\} \\ W &= \left\{ e^{i\theta} \mid \theta \in \left(\frac{\pi}{2} - \epsilon, \frac{\pi}{2} + \epsilon \right) \right\} \\ E &= \left\{ e^{i\theta} \mid \theta \in \left(-\frac{\pi}{2} - \epsilon, -\frac{\pi}{2} + \epsilon \right) \right\}, \end{aligned}$$

where $\epsilon \in (0, \pi/2)$ is a number (in some sense characterizing the quality of the observations being made: a smaller ϵ means better precision). The observer may, initially be unaware or undecided regarding the value of ϵ , so that two types of situations may occur: one with $\epsilon \leq \pi/4$ and the other with $\epsilon > \pi/4$ see figure 1. We set

$$H = \{N, S, W, E, N^c, S^c, W^c, E^c\}$$

and ask the reader to verify the pictures of the corresponding graphs Γ_H presented in figure 2. Observe how different the two graphs are. Does one of them resemble X more than the other? Oddly enough, it is the *lower* quality observation that provided the better picture. How come?

The preceding example illustrates an idea central to this paper: memory is (or at least *should be*) a geometric/combinatorial model (here, a graph) of the state space X , marked with additional information (content at the vertices). It seems reasonable to also mark a point of this geometric model (in Γ_H that would mean a vertex, or in other words, an H -visible state) to represent the observer's conjecture about the current state of X .

The categorical point of view discussed earlier will now work as follows: adding sensors to \mathcal{O} corresponds to an inclusion of the list H in a longer list \tilde{H} , creating a refinement $\mathcal{P}(\tilde{H})$ of the partition $\mathcal{P}(H)$ and hence a map of the graph $\Gamma_{\tilde{H}}$ onto the graph Γ_H induced by sending every \tilde{H} -visible state to the unique H -visible state containing it. Thus the category \mathcal{C} for our modeling problem may be taken to have weighted connected bipartite graphs, and an arrow from an object Γ to an object Γ' is a surjective graph morphism (edges may be contracted) from Γ' onto Γ .

A major problem standing in the way of developing this approach further is that Γ_H can be quite arbitrary on the large scale, so that searching Γ_H may turn out to be computationally unfeasible as the observer \mathcal{O} comes to possess more and more sensors and H grows in size accordingly.

The situation is even worse than that: we have just assumed too much. First of all, the angel's job is *only* to do the book-keeping for \mathcal{O} . In our interpretation of memory, this means the angel has no business representing the observer's questions by objectively defined complementary pairs of subsets of X : \mathcal{O} will most probably *never* have that much information about *any* of its sensors. This angel has to go home then, and we need to face the fact that the observer has to maintain Γ_H and its content on his own, without any prior knowledge about the structure of H and without any assurance that the content of the vertices of Γ_H is objectively correct: \mathcal{O} is only able to *sample* states from X by making repeated observations, so the probabilities recorded at the vertices may be very different from the objective ones.

To deliver the final blow, consider this: even if there is a way for \mathcal{O} to magically keep track of the objective structure of Γ_H , the observer may still be required to restructure Γ_H numerous times as time evolves. For example, adding a vertex to Γ_H where he initially thought there was none: suppose that, up to time $t = t_0$, \mathcal{O} has never sampled a state for which both questions a and b had a positive answer; as a result, up until time $t = t_0$, no vertex in Γ_H listing a positive answer to a lists a positive answer to b ; if at time t_0 the observer suddenly makes the observation that $x \in a \cap b$ for some state $x \in X$, then Γ_H must be updated to reflect that observation. Can this be done quickly and efficiently?

The answer is definitely negative: example 1.1 and figure 2 demonstrate the fact that Γ_H can change immensely as a result of seemingly minor adjustments of H (e.g., in the example, if ϵ is very close to $\pi/4$, a very small change in ϵ will cause a cycle to collapse into a tree or vice-versa). Given that Γ_H may resemble practically any big graph on the large scale, the problem of structural updating for Γ_H becomes intractable very quickly.

Despite the failure of our first attempt, we will construct a workable model using an almost identical skeleton of ideas. After all, Γ_H does have some very desirable traits one would be happy to retain:

- Γ_H encodes rudimentary logic (negation, implication),
- it seems possible to synchronize updating content with structural updating (e.g., erase vertices representing visible states of little interest),
- a *learning goal* for \mathcal{O} can be defined: have Γ_H grow to be big/detailed enough, with its content eventually close enough to objective values, in order to guarantee the success of \mathcal{O} .

1.3. Results. The model of memory we propose in this paper, has an underlying category whose objects are duals, in some sense, of graphs belonging to the well studied family of *median graphs*. The interested reader should see [Rol98] for an extended bibliography on the subject and a detailed treatment of duality theory for median algebras, and median graphs in particular. A very good treatment of median graphs and their duality theory is given in [Nic04]. We give a self-contained exposition of the relevant notions and results in section 2, though we do focus on the current application and omit proofs of results from the literature.

One remarkable feature of our model is that the structure of a median graph itself provides our observers with a rudimentary sense of logic. We will discuss this formal idea of ‘common sense’ as we progress through the next section.

Searching our databases for content and updating content is discussed in section 3. We explain how searching our database structures essentially coincides with content updating and discuss some of the implications regarding learning. A possible implementation of the model as a network of neuron-like elements is offered as well, and we discuss its computational efficiency.

Section 4 is dedicated to structural updating and a discussion of how the phenomena of learning, language formation, forgetting and understanding are realized in the model.

As the exposition evolves, we periodically pause to look at how various natural phenomena are accounted for by our model. Right now we can state, with some satisfaction, that our model explains an overwhelming majority of the phenomena we had already listed as curious aspects of the human thought process and human memory.

The last section discusses weaknesses of the model and possible ways to get rid of them – a topic for future research.

Acknowledgements. The author is deeply indebted to Michael Jablonski and Vera Tonic for proof-reading the text and for numerous useful comments on both content and form. Many thanks to Lucas Sabalka for the idea of replacing computationally cumbersome examples of spaces with walls with the compass example (which is yet to reappear in our narrative); to Krešimir Josić for commentary on the material of section 3 and ongoing discourse. The author is a newcomer to the field, and must confess limited knowledge of existing literature. I am grateful in advance for any comments and criticisms, and will happily acknowledge credit wherever credit is due.

2. MODELING MEMORY USING POC-SETS AND MEDIAN GRAPHS

2.1. Poc-sets. We consider an abstract version of a system of questions, due to Roller [Rol98]:

Definition 2.1 (poc-set). A poc-set $(P, \leq, *)$ is a partially-ordered set (P, \leq) with a minimum (denoted by 0), endowed with an order-reversing involution $a \mapsto a^*$ such that $a \leq a^*$ implies $a = 0$ for all $a \in P$. The maximum $0^* \in P$ is denoted by 1; the elements $0, 1 \in P$ are said to be *trivial*; all other $a \in P$ are *proper*.

Example 2.2 (standard poc-sets). The Borel σ -algebra \mathcal{B} carries a natural poc-set structure $(\mathcal{B}, \subseteq, a \mapsto a^c)$. If P is a poc-set and $0 \in Q \subseteq P$ satisfies $Q^* = Q$ then Q is a (sub) poc-set (of P).

For any proper $a, b \in P$, it is easy to see that only one of the following may hold:

$$(1) \quad a \leq b, \quad a \leq b^*, \quad a^* \leq b, \quad a^* \leq b^*.$$

Definition 2.3 (nesting, transversality). A pair a, b of elements in a poc-set P is said to be *nested*, if any one of the relations in 1 holds. Otherwise, the pair a, b is said to be *transverse* (denoted $a \pitchfork b$). A subset $S \subseteq P$ is said to be nested (resp. transverse), if the elements of S are pairwise nested (resp. transverse).

A means for relating poc-sets to each other will be required:

Definition 2.4 (morphisms). A function $f : P \rightarrow Q$ between poc-sets is a *morphism* if $f(0) = 0$, f is order-preserving and $f(a^*) = f(a)^*$ for all $a \in P$.

We were previously considering the possibility of modeling the memory of an observer \mathcal{O} by a sub poc-set of \mathcal{B} , but it works better to use a pair (P, f) , with P an abstract poc-set and $f : P \rightarrow \mathcal{B}$ a morphism. We shall presently see that the abstract object P gives rise to a graph $\Gamma(P)$ such that: (a) P is reconstructible from $\Gamma(P)$ and (b) the graph $\Gamma(P, f) = \Gamma_{f(P)}$ of visible states determined by $f(P)$ (see discussion in sub-section 1.2) is canonically embedded in $\Gamma(P)$. Thus, if we choose $\Gamma(P)$ to represent the memory of \mathcal{O} while f is viewed as an interpretation of P in the reality presented by X , then the above property of $\Gamma(P)$ implies that no updating of the structure of the memory graph $\Gamma(P)$ is required so long as the implication relations among the questions available to \mathcal{O} (that is – the structure of P) remain unchanged. Not taking other aspects of the modeling problem into account, properties (a) and (b) of $\Gamma(P)$ should be viewed as the main argument in favor of preferring $\Gamma(P)$ over $\Gamma(P, f)$: while containing complete information about P (and hence not being ‘too big’), $\Gamma(P)$ has ‘sufficient space’ to accommodate *any* possible interpretation of P in X .

Perhaps it is time again for a concrete example.

Example 2.5 (Compass, part 2). We return to example 1.1. Recall the sets N, S, W, E defined there as arcs of length 2ϵ , $\epsilon \in (0, \pi/2)$, on the unit circle – see figure 3 and example 1.1. This time, define a formal poc-set P as follows:

$$P = \{0, 1, n, s, w, e, n^*, s^*, w^*, e^*\},$$

and we set $f(n) = N$, $f(s) = S$ and so on, together with the relations following from the requirement that f be a morphism.

We will compare two situations, see figure 3: (i) the one where $\epsilon \in (\pi/4, \pi/2)$ with (ii) the one where $\epsilon \in (0, \pi/4)$.

Observe that $\epsilon < \pi/2$ implies that a positive answer to n excludes a positive answer to s and vice versa. The same is true about the pair $\{e, w\}$. We have chosen ϵ in this way in order to enable a choice of the poc-set structure on P having the relations $n \leq s^*$ and $w \leq e^*$ and all conclusions thereof. Both in this example and in its continuation (example 2.9) P will be chosen to have these relations (and their consequences). To be




FIGURE 3. Observing a compass (1). Two possible realizations of the same abstract poc-set structure.

sure, we list the relations on P :

$$\begin{aligned} n &< s^*, & n \pitchfork e, & n \pitchfork w, \\ s &< n^*, & s \pitchfork e, & s \pitchfork w, \\ e &< w^*, & w &< e^*. \end{aligned}$$

Note that if we had wanted to allow similar realizations of P with $\epsilon > \pi/2$, we would have been forced to change all the above relations into transversality relations, in which situation the poc-set P on its own would not have held *any* information about the observed system.

When $\epsilon < \pi/4$ (see (ii) in the figure), the sets N, S, E, W are pairwise complementary. This means that the poc-set $f(P)$ has the relations $N < E^c$ and $N < W^c$ although neither $n < e^*$ nor $n < w^*$ hold in P .

For $\epsilon > \pi/4$ though (see (i) in the figure), all the relations holding in $f(P)$ are accounted for already in P . In example 2.9 we will see how this difference between the situations is visualized at the level of the corresponding graphs.

2.2. Sageev-Roller duality. The construction of $\Gamma(P)$ – the graph *dual* to P – given in this paragraph is originally due to Sageev [Sag95] in a special case. The presentation we have chosen and the discussion of dual maps is due to Roller [Rol98]. However, we have chosen to alter Roller’s original terminology to better fit the intended application.

Definition 2.6 (coherence, vertices). Let $(P, \leq, *)$ be a poc-set. A subset $\alpha \subset P$ is said to be coherent if $a \leq b^*$ holds for *no* $a, b \in \alpha$. A maximal coherent family will be called a *vertex* of the graph $\Gamma(P)$.

It is easy to see that a coherent family $\alpha \subset P$ lies in $V\Gamma(P)$ iff α is a $*$ -selection on P (meaning that for all $a \in P$, either $a \in \alpha$ or $a^* \in \alpha$, but not both).

Example 2.7. Let $\mathbf{2} = \{0, 1\}$ be the *trivial poc-set*, with the obvious relation $0 < 1$. Show that the assignment $f \mapsto f^{-1}(1)$ from the set $\text{Hom}(P, \mathbf{2})$ of all poc-morphisms of $f : P \rightarrow \mathbf{2}$ to the set $V\Gamma(P)$ is a bijection.

Definition 2.8 (edges). Let $(P, \leq, *)$ be a poc-set and $u, v \in V\Gamma(P)$. We set $\{u, v\} \in E\Gamma(P)$ iff $u \Delta v = \{a, a^*\}$ for some $a \in P$. Here $u \Delta v$ denotes the symmetric difference $(u \setminus v) \cup (v \setminus u)$.

More generally, for $u, v \in V\Gamma(P)$ one has $|u \Delta v| = 2|u \cap v^*|$, which implies that the expression $\Delta(u, v) = \frac{1}{2}|u \Delta v|$ – the number of questions separating u from v – is a distance function on $V\Gamma(P)$. Moreover, $\Delta(u, v) = 1$ iff u and v are joined by an edge in $\Gamma(P)$.

Example 2.9 (Compass, part 3). We would like to go back to the poc-set

$$P = \{0, 1, n, s, w, e, n^*, s^*, w^*, e^*\}$$

defined in example 2.5 and endowed with the structure

$$\begin{aligned} n < s^*, & \quad n \pitchfork e, & \quad n \pitchfork w, \\ s < n^*, & \quad s \pitchfork e, & \quad s \pitchfork w, \\ e < w^*, & \quad w < e^*. \end{aligned}$$

In order to visualize $V\Gamma(P)$, consider $\Gamma(P)$ as a full simple subgraph of the cube with vertex set 2^P (where each vertex corresponds to a choice of one answer for every question, ignoring coherence issues): $\Gamma(P)$ is the result of erasing all vertices of 2^P turning out to be incoherent, together with their adjacent edges – see figure 4.

Example 2.10 (The n -cube). Let T be a set of n distinct symbols, and let P be the poc-set generated by T with no relations, that is: $P = \{0, 0^*\} \cup T \cup T^*$, where T^* is the set of symbols of the form t^* such that $T^* \cap T = \emptyset$, and no two distinct proper elements of P are comparable.

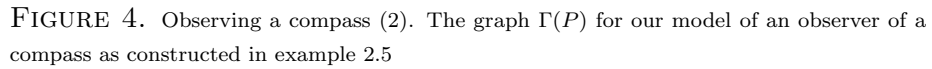
Then any $*$ -selection on P is a maximal coherent subfamily, and we conclude that $\Gamma(P)$ is the 1-dimensional skeleton of the n -dimensional cube.

The next example is based on the following recommended easy exercise:

Example 2.11. We offer the following exercise to the reader. Suppose P is a finite poc-set and let $a \in P$. Prove that the following are equivalent:

- (1) a is nested with every element of P ,
- (2) $\Gamma(P)$ has one and only one edge $\{u, v\}$ satisfying $u \Delta v = \{a, a^*\}$, and this edge is a cut-edge.

Example 2.12 (Trees and n -pompoms). The above exercise shows that if P is nested then $\Gamma(P)$ is a tree. The converse is known to be true as well (exercise). A special case that is very easy to compute is constructed as follows: take A to be a set of n distinct symbols and let $P = \{0, 0^*\} \cup A \cup A^*$ where A and A^* are disjoint and subject to the relations $a_i^* < a_j$ for all $1 \leq i < j \leq n$. This immediately implies there are only two kinds of vertices: the vertex $v_0 = \{0^*\} \cup A$ and the vertices $v_i = \{0^*, a_i^*\} \cup (A \setminus \{a_i\})$. Then $\Gamma(P)$ is the tree with $(n + 1)$ vertices and n leaves v_1, \dots, v_n . We will refer to this structure as the n -pompom.


$$(2) \quad I(u, v) = \{w \in V\Gamma \mid d_\Gamma(u, w) + d_\Gamma(w, v) = d_\Gamma(u, v)\},$$

Definition 2.13 (median graph). A connected graph Γ is a median graph if, for all $u, v, w \in V\Gamma$ the intersection $I(u, v) \cap I(v, w) \cap I(u, w)$ contains precisely one vertex (denoted $\text{med}(u, v, w)$).

An example visualizing what is going on is the standard rectangular integer grid:




FIGURE 5. The Integer Grid: example of computing a median in a median graph.

Example 2.14 (The Grid). Let X be the set of points $(x, y) \in \mathbb{R}^2$ satisfying $x \in [0, m+1]$ and $y \in [0, n+1]$ with m, n being positive integers. Let

$$h_s = \left\{ (x, y) \mid y < s + \frac{1}{2} \right\}, \quad v_t = \left\{ (x, y) \mid x < t + \frac{1}{2} \right\},$$

where $s \in \{1, \dots, n\}$ and $t \in \{1, \dots, m\}$. Then the set

$$P = \{\emptyset, X\} \cup \{h_s, h_s^c\}_{s=1}^n \cup \{v_t, v_t^c\}_{t=1}^m$$


is a sub poc-set of \mathcal{B} (the Borel σ -algebra on X) with respect to inclusion (\subseteq) and complementation ($A \mapsto A^c$). It is easy to see that $h_s \cap v_t$ for all s, t , while $h_s \subset h_{s+1}$ and $v_t \subset v_{t+1}$ for all relevant s and t . As a result, any subset of P is consistent if and only if it is coherent, and the visible states of X with respect to P (letting P be realized by the inclusion map) are in one-to-one correspondence with the vertices of $\Gamma(P)$. Observe that the visible states are in one-to-one correspondence with the integer points of X , so it makes sense to use these points as representatives, joining two such points by a straight line segment if and only if the corresponding vertices of $\Gamma(P)$ are joined by an edge. Figure 5 demonstrates the computation of a median of three vertices of $\Gamma(P)$, drawn over a diagram of $\Gamma(P)$ realized in this way in the plane.

A good exercise for the interested reader will be to prove the following –

Lemma 2.15 (median operation). If $(P, \leq, *)$ is a finite poc-set, then $\Gamma(P)$ is a median graph, and $\text{med}(u, v, w) = (u \cap v) \cup (v \cap w) \cup (u \cap w)$ for all $u, v, w \in V\Gamma(P)$.

Theorem 2.16 (‘Sageev-Roller duality’, [Rol98]). If P is a finite poc-set then $\Gamma(P)$ is a connected median graph and $d_{\Gamma(P)}$ coincides with Δ . Furthermore, every finite median graph arises in this way.

By this theorem, the construction of an arbitrary ‘memory graph’ makes it automatically a median graph. Also, every median graph can be thought of as a ‘memory graph’

FIGURE 6. Visualizing the dual map of f from example 2.19.

when provided with a realization map relating it to a state space of an observed system. We conclude that the property of being a median graph can be identified as a one of the “guiding principles” mentioned in the introduction – a principle restricting the structure of such graphs.

We shall now proceed to demonstrate the applications of this idea using the study of morphisms of median graphs (maps between median graphs preserving the median structure).

An additional aspect of the above duality is that morphisms of poc-sets translate into median morphisms of median graphs and vice versa. If A and B are median graphs, then a morphism from A to B is a function $VA \rightarrow VB$ preserving medians.

Remark 2.17. We do not require a morphism of median graphs to preserve the adjacency relation. In fact, imposing this additional requirement proves to be overly restrictive for our purposes.

If $f : P \rightarrow Q$ is a morphism of finite poc-sets, then a *dual morphism* of median graphs $f^\circ : \Gamma(Q) \rightarrow \Gamma(P)$ is defined by $f^\circ(v) = f^{-1}(v)$ for $v \in V\Gamma(Q)$. Some properties of this construction are:

Proposition 2.18 (functorial properties, see [Rol98]). If $f : P \rightarrow Q$, $g : Q \rightarrow R$ are morphisms of finite poc-sets then:

- (1) $(g \circ f)^\circ = f^\circ \circ g^\circ$;
- (2) f is surjective if and only if f° is injective;
- (3) f is an embedding if and only if f° is surjective.

(by an embedding we mean an isomorphism onto the image)

Example 2.19 (see fig. 6). Let P and Q be poc-sets with 4 proper elements each – a, a^*, b, b^* – such that $a \pitchfork b$ in P but $a < b$ in Q . In this case, the set-theoretic

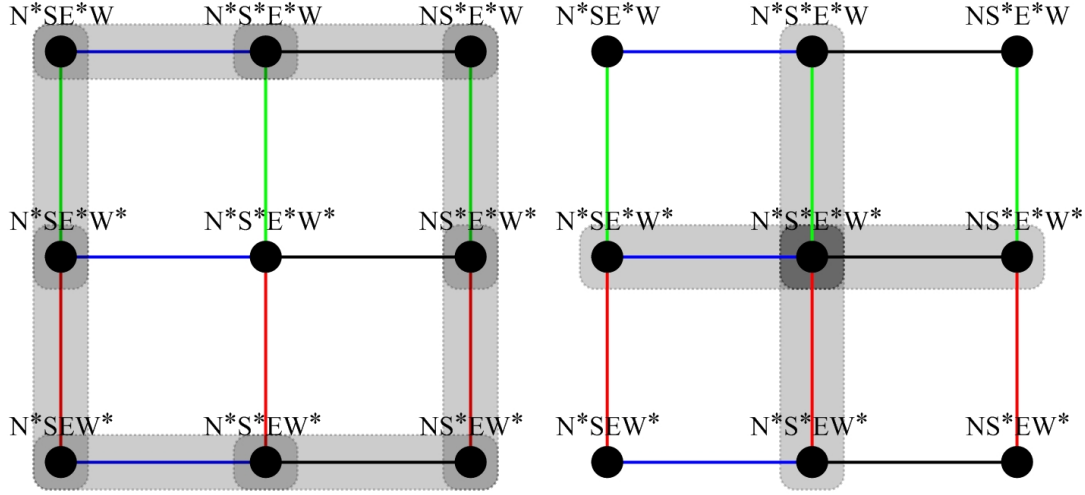


FIGURE 7. Observing a compass (2). The two realizations from example 2.5 induce *very* different embeddings of $\Gamma(P, f)$ in $\Gamma(P)$: for $\epsilon > \pi/4$ (left) and for $\epsilon < \pi/4$ (right). Edges of $\Gamma(P)$ belonging to $\Gamma(P, f)$ are marked by transparent boxes drawn on top of them.

identity map $f = id : P \rightarrow Q$ is a morphism, while its inverse is not. $\Gamma(P)$ is a 4-cycle, while $\Gamma(Q)$ is a path of length 2 (having 3 vertices). While f is bijective, it is not an embedding: though f^{-1} is well-defined, it is not a morphism of poc-sets.

2.3. Consistent families and weights: our model of an observer. We now return to the idea of a morphism $f : P \rightarrow \mathcal{B}$ representing the memory of an observer at a fixed moment in time. We will henceforth refer to such f as a *representation* of P .

For every $x \in X$, consider the maximal coherent subfamily of \mathcal{B} :

$$(3) \quad \pi_x = \{A \in \mathcal{B} \mid x \in A\}.$$

One can apply f° to π_x to obtain the maximal coherent family

$$(4) \quad \pi_{P,f}(x) = \{a \in P \mid x \in f(a)\} \in V\Gamma(P).$$

We now have a map $\pi_{P,f} : X \rightarrow V\Gamma(P)$ selecting precisely those vertices in $\Gamma(P)$ corresponding to visible states of X relative to $f(P)$. Since $\pi_{P,f}$ is constant on every visible state, $\pi_{P,f}$ induces an injective map from $V\Gamma(P, f)$ into $V\Gamma(P)$. From the definition of edges in both graphs it is clear that this injection is an embedding of graphs, which we denote by $\bar{\pi}_{P,f} : \Gamma(P, f) \rightarrow \Gamma(P)$.

Definition 2.20 (consistent families). A set $u \subseteq P$ is said to be consistent relative to a representation f , if there is a point $x \in X$ such that $f(u) \subseteq \pi_x$.

In particular, a vertex u of $\Gamma(P)$ is consistent iff it lies in the image of $\Gamma(P, f)$ under $\bar{\pi}_{P,f}$. We illustrate this on the compass example (examples 2.5 and 2.9):

Example 2.21 (Compass, part 4). Figure 7 shows $\Gamma(P, f)$ embedded in $\Gamma(P)$. It is important to observe that the central vertex is inconsistent for $\epsilon > \pi/4$, whereas the

same vertex becomes consistent upon reducing ϵ to a value less than $\pi/4$, with the corner vertices becoming inconsistent in this situation.

Looking at the picture leaves one amused by the fact that, somehow, smaller precision in answering the questions n, s, w and e (corresponding to $\epsilon > \pi/4$) contributes to a better discrete model of X (recall X was a circle) than the one obtained from more precise answers. More than that, we consider this example as clear proof of our claim that separating the logical structure P from its realization (resulting in making the distinction between $\Gamma(P)$ and $\Gamma(P, f)$) constitutes a significant improvement over the simplistic model described in the introduction: indeed, $\Gamma(P)$ in this case is spacious enough to include faithful representations of both types of realization that we had discussed, while the separate realizations seem completely incompatible (compare the two versions of $\Gamma(P, f)$ in figure 7 again).

Now we are ready to formally define a model of an observer:

Definition 2.22. Let (X, \mathcal{B}, μ) be a probability space. An *observer* of X is a quadruple $\mathcal{O} = (P, f, p, \epsilon)$, where P is a finite poc-set, $f : P \rightarrow \mathcal{B}$ is a morphism of poc-sets, $p : VT(P) \rightarrow [0, \infty)$ is a function on the vertices of $\Gamma(P)$ and ϵ is a subset of P .

The additional data elements – p and ϵ – are new to our discussion. Ethological¹ considerations seem to imply that memory is subject to prioritization: different events are considered relevant to different extents depending on the type of conflict they create between the individual and the environment. The types of conflict can be classified, and the subjective perception of the intensity of an interaction between the individual and the environment may be measured by observing the change in the level of physiological stress resulting from the interaction [Gur]. In this context, we want $p(v)$ to represent an estimation that \mathcal{O} has regarding the relevance of the event $\bigcap_{a \in v} f(a)$, for every $v \in VT(P)$, and ϵ to represent the conjecture \mathcal{O} has regarding the current state of the universe. Ideally, ϵ will be a vertex of $\Gamma(P)$. Less ideally, ϵ is a coherent family. One should note, however, that humans holding a completely coherent set of convictions are rarely to be found.

We will need some new notation:

Definition 2.23 (convexity, convex hull). In a graph Γ , one says that a subset W of VT is *convex*, if $I(u, v)$ is contained in W for all $u, v \in W$. The *convex hull* $conv(W)$ of W is defined to be the intersection of all convex subsets of VT containing W .

Clearly, $conv(W)$ is the smallest convex subset of VT containing W .

Definition 2.24. Let P be a poc-set. For $a \in P$ and $A \subseteq P$ we denote:

$$(5) \quad V(a) = \{u \in VT(P) \mid a \in u\}, \quad V(A) = \bigcap_{a \in A} V(a).$$

Observe that $V(a^*) = VT(P) \setminus V(a)$ for all $a \in P$. Thus, the family $\{V(a)\}_{a \in P}$ forms a poc-set with respect to inclusion and complementation. In fact, it is easy to see that this poc-set is isomorphic to P via the map $a \mapsto V(a)$. For the case of a median graph,

¹Ethology – a field of Biology, studying animal behaviour.

when one can safely write $\Gamma = \Gamma(P)$, it turns out that $\text{conv}(W) = \bigcap_{W \subset V(a)} V(a)$. Thus, $W \subset V\Gamma(P)$ is convex if and only if it equals the intersection of a family of halfspaces of $V\Gamma(P)$.

Let us now return to the discussion of the relationship between a living observer and our notion of an observer $\mathcal{O} = (P, f, p, \epsilon)$. We first focus our attention on the function p . As stated above, p should be thought of as a measure of the relevance of a vertex $u \in V\Gamma(P)$ in the observer's eyes: the event u is likely to be disregarded by the observer if $p(u)$ is negligible, while an event of the form $V(A)$ will attract more of the observer's attention the higher the cumulative value $p(V(A)) = \sum_{u \in V(A)} p(u)$. Thus, from an information-theoretic point of view, every event $F \subset V\Gamma(P)$ has a probability attached to it:

$$(6) \quad \Pr_{\mathcal{O}}(F) = \frac{p(F)}{p(V\Gamma(P))}$$

This should not be confused with the function taken on by the set ϵ corresponding to the observer's conjecture of the 'current state of affairs': $\Pr_{\mathcal{O}}(V(\epsilon))$ being small means our observer is aware of no sources of stress right now; the same quantity being large *may* provide a motivation for our observer to take action to change ϵ in the direction of lowering the stressfulness of the situation.

However, one should not imagine living organisms as trying to solve some kind optimization problem (with the objective being to minimize stress), but, rather, as trying to solve an equilibrium problem. Ethological studies [Gur] show that minimization of stress could not be regarded as a plausible goal for *every* observer. Our information-theoretic interpretation of p is a convenient tool for describing this phenomenon: an observer cannot be attracted to the idea of pushing X into a state which this same observer sees as uninteresting.

Here are some possible combinations of f , ϵ and p to contemplate:

Example 2.25 (objectivity). Given $\mathcal{O} = (P, f, p, \epsilon)$, choose $p(u) = \mu(\bigcap_{a \in u} f(a))$ and $\epsilon = \pi_x$ where x is the actual current state of the system. The resulting observer is, in some sense, objective: his perception of events is precise (ϵ is a maximal coherent family), up to date ($V(\epsilon)$ includes the current state) and unbiased ($\Pr_{\mathcal{O}}((\cdot))$ gives the correct probabilities of events).

Example 2.26 (misperception). Once again, suppose P and f are given, and x is the current state of the universe. Consider some possible situations for an observer $\mathcal{O} = (P, f, p, \epsilon)$:

- $\epsilon \subsetneq \pi_{P,f}(x)$: This means that \mathcal{O} has a question allowing it to improve its perception of the current state of the universe.
- $\epsilon \not\subseteq \pi_{P,f}(x)$: This means that \mathcal{O} has a question which – had it been asked – would have proved the observer's view of the universe to be inconsistent with the current state, forcing the observer to alter its structure as an observer.
- $p(u) = 0$ **with u consistent**: In this situation \mathcal{O} regards an existing state of the universe as impossible/negligible.

$p(u) > 0$ **with u inconsistent:** Here an impossible event (u) may attract the attention of \mathcal{O} , who considers this event as disturbing (and hence probable).

The above situations demonstrate the flexibility of our model. They also motivate looking for a way to measure how far an observer \mathcal{O} is from being objective.

2.4. $\Gamma(P)$ as a representation of ‘common sense’. The preceding examples are concerned with the possible relations between the content components p and ϵ of an observer $\mathcal{O} = (P, f, p, \epsilon)$ and its representation map. Now we would like to focus on the structural component P and its relation to so-called ‘common sense’.

By common sense we do not mean ideas or rules common in society: ‘cannibalism is bad’ is not an inborn common notion, but a fundamental and non-trivial cornerstone of modern society. By common sense we mean a certain shared notion of rudimentary logic that is inherent to all our actions. In part, this notion is embodied in our model through the assumption that memory is structured by a poc-set, and realized by a poc-morphism. Consider an observer \mathcal{O} as in definition 2.22. Both the structure of P and that of $\Gamma(P)$ carry information about implication relations among events in X as those are perceived by \mathcal{O} . This is due to the equivalence $a < b \Leftrightarrow V(a) \subset V(b)$ holding for all $a, b \in P$. Together with the fact that $V(a^*) = V(a)^c$, this provides an automated tool \mathcal{O} can use for ‘sub-conscious’ reasoning.

Thus, all our observers share the way in which the content of their memory is ordered, and this fact is bound to affect the ways in which two observers sharing the same environment (or territory) synchronize their actions.

When is it that we are able to demonstrate to others that we understand something well? The pedagogical answer has always been that good understanding is defined as a state in which discussing the subject matter adequately from a logical standpoint does not require a lot of conscious effort.

However, logic as we understand it involves the ability to operate with combined observations. While implication and negation are inherent to observers through the poc-set structure of the ‘atomic’ statements, conjunctions (and, dually, disjunctions) are not taken into account directly by the poc-set structure.

To demonstrate this, fix a positive integer n and consider two hypothetical observers:

- \mathcal{O}_1 with poc-set component P_1 such that $\Gamma(P_1)$ is the n -cube (see example 2.10), and an excitation function p_1 assigning a unit value to every vertex;
- \mathcal{O}_2 with poc-set structure P_2 such that $\Gamma(P_2)$ is the 2^n -pom-pom (see example 2.12) with center v_0 , and with p_2 assigning unit excitation values to all the leaves of $\Gamma(P_2)$ and null excitation to v_0 .

Furthermore, suppose f_1 realizes \mathcal{O}_1 so that the partition $\mathcal{P}(f_1(P_1))$ induced on X is uniform with respect to μ (in particular, \mathcal{O}_1 is objective). We do not want \mathcal{O}_2 to be at an unfair disadvantage, so assume the partition $\mathcal{P}(f_2(P_2))$ coincides with $\mathcal{P}(f_1(P_1))$ so that v_0 is f_2 -inconsistent. Thus, \mathcal{O}_2 is objective as well, and the only way in which the two observers differ is the combinatorics each of them uses to model the same partition of X .

We will now ask the standard question an information theorist asks in situations like this: what is the expected minimum number of observations that each observer needs to make about the current state in order to identify the vertex in its memory corresponding to this state of X ? Then answers are clearly very different: n for \mathcal{O}_1 and $2^n - 1$ for \mathcal{O}_2 .

A natural way for \mathcal{O}_2 to come closer to the optimal position of \mathcal{O}_1 is to widen the supply of direct observations \mathcal{O}_2 is able to make. If \mathcal{O}_2 had a mechanism for expanding P_2 by adding conjunctions of elements from P_2 to it, the updated version of \mathcal{O}_2 would be able to perceive X with lower entropy.

On the other hand, it is not reasonable for \mathcal{O}_1 to maintain the strategy of keeping P_1 completely transverse (except for trivial nesting relations) over a long time: as the number of available observation tools increases, $\Gamma(P_1)$ will keep growing exponentially, without ever encoding *any* of the recorded information in its combinatorial structure.

The inevitable conclusion is that an approach balancing content and structure should exist, and we expect it to become all the more efficient as the particular implementation of the memory structure comes to possess tools allowing the observer to refine its observations by combining them at will.

3. POSSIBLE IMPLEMENTATION AND INFORMATION RETRIEVAL

3.1. The basic searching problem. The goal of this section will be to define the basic searching problem for an observer $\mathcal{O} = (P, f, p, \epsilon)$ and to discuss the main features of an implementation that one may regard as efficient.

Deferring updating tasks which involve altering the structure of P or the excitation function p , one is left with the task of efficiently updating ϵ – the observer’s conjecture about the current state of events – in response to an incoming observation.

The basic search/retrieval/update problem may be formulated as follows. Suppose an observer modeled by $\mathcal{O} = (P, f, p, \epsilon)$ has just made the observation $a \in P$. Then there is a need to (1) decide whether $\epsilon \cup \{a\}$ is consistent, and then (2) replace ϵ by $\epsilon \cup \{a\}$ in case it is, or (3) replace ϵ by a new description of the perceived current state that is as consistent as possible with ϵ and the new observation.

This is the point when any discussion of efficiency will depend on the specific implementation, that is: on *how* the memory structure is realized and maintained. In order to facilitate this discussion, let us consider the different components of an observer \mathcal{O} from a more practical point of view:

- Our initial assumption was that observing (and recording observations about) a system with a given state space X may be thought of as maintaining a database whose purpose is to mimic the geometry and topology of X , as those are revealed through asking binary questions about X . Such a question can be very basic, e.g. “Is neuron number 1234567 firing right now?”.

- The function f in the definition of an observer provides the connection to reality: f essentially represents the sensors used by the observer for watching the evolution of the observed system, e.g. neuron 1234567 fires if and only if a certain patch of light receptors on our retina gets hit by a sufficient number of photons.
- Both $\Gamma(P)$ and P are maintained (stored in memory) as directed graphs, with the vertices of $\Gamma(P)$ labeled by their values under the function p and the edges of $\Gamma(P)$ labeled by elements of P .
- Finally, recording ϵ is no more than a labeling on the vertices of the graph representing P . This graph is nothing but a Hasse diagram with additional edges labeled by $(*)$ to join every $a \in P$ to $a^* \in P$. The labeling corresponding to ϵ works as follows: a vertex $a \in P$ is ‘ON’ if and only if $a \in \epsilon$.

3.2. Idealized searching. Let us now suppose that an implementation of P is equipped with the following *idealized* features in addition to the graph structure we had just described:

Propagating Excitation: Every node of P (we identify elements of P with the corresponding nodes) has an excited state and a non-excited state. If $a \in P$ is excited, then so is every $b \in P$ with $b > a$.

Contradiction Detection: If $a \in P$ is excited, and a^* is ‘ON’, then a flag is raised and the detector outputs a .

The first feature is motivated by neurons, but very far from being a precise copy of the same mechanism. Upon observing the algorithmic implications of the idealized feature we will de-idealize it and discuss the consequences.

Suppose an observer modeled by \mathcal{O} makes an observation $a \in P$. The situation then requires a reaction on the part of the observer with the aim to keep \mathcal{O} up to date: (1) if $a \in \epsilon$ then ϵ is not changed; if $a \notin \epsilon$, then either (2) $\epsilon \cup \{a\}$ is coherent and ϵ will be replaced by $\epsilon \cup \{a\}$, or (3) $\epsilon \cup \{a\}$ is incoherent and ϵ must be replaced by a ‘closest approximation’ ϵ' containing a .

The entire process begins with switching a to an excited state. The excitation propagates along P . Recall that for a coherent ϵ , $\epsilon \cup \{a\}$ is incoherent iff there exists $b \in \epsilon$ such that $a \leq b^*$. Thus, under the operative assumption that ϵ is coherent (which may well be false), $\epsilon \cup \{a\}$ is proved to be coherent iff the contradiction detector raises no flags as a result of our exciting a . In this case ϵ will be replaced by $\epsilon \cup \{a\}$ – turn a on if it was off – and we are done. In the situation when a contradiction is detected at $b \in P$ ($b \geq a$), simply turn b^* off and turn b on for all such b in order to obtain ϵ' .

To gain a better understanding of the actual meaning of this updating process, we examine it for the special case when ϵ is a vertex of $\Gamma(P)$ (that is: \mathcal{O} is completely decided – though not necessarily right – about the current state of the observed system). In this case it is easy to see that ϵ' is the unique vertex satisfying $\epsilon' \in I(\epsilon, v)$ for all $v \in V(a)$. Thus, in this case ϵ' is, in a sense, the best possible approximation of ϵ by elements of $V(a)$. Intuitively, we think there is no better candidate for ϵ' .

Now, it is unrealistic to assume that the propagation process takes no time, while it is reasonable to assume that the cells of the realization of P (the nodes of the graph) have equal physical characteristics. Thus, the excitation signal will propagate through P at a linear pace, implying our updating algorithm produces ϵ' in a time linear in the length of the longest maximal chain (in P) joining a with b , where $b > a$ and $b^* \in \epsilon$. Clearly, this is as efficient as one may hope for. Of course, the propagation property plays the role of an extremely powerful parallel processor with the capacity to handle a potentially unbounded number of parallel computation threads.

3.3. More realistic searching. The reader will have noticed by now the similarities between our idea of propagation of excitation and the propagation of signals in a network of neurons. Consequently, the reader will have thought of the excessive optimism embodied in the assumption that the propagated signal does not dissipate: a neuron ν fires upon accumulating sufficient charge on its dendrites, and this charge is then distributed along the axons into the synapses; the more neurons have their dendrites connected to the axon of ν , the less charge will accumulate on each of them. It is therefore reasonable to expect the signal to dissipate exponentially fast, if P is sufficiently branched.

What does this mean for our updating algorithm?

First of all, it becomes possible for some of the elements $b \in \epsilon$ which satisfy $a < b^*$ to remain in ϵ' when the propagating excitation wave emanating from a does not reach the corresponding nodes, thus failing to trigger the contradiction detectors.

As a result, ϵ' will not be coherent. This is the reason why we did not require coherence from ϵ in the definition of an observer, as well as our reason for emphasizing the possibility of ϵ being incoherent throughout the preceding discussion. Perhaps this is also the reason why we rarely observe humans with a completely coherent view of the world.

One should not completely despair of the task of ultimately updating ϵ with all the relevant elements of P , though. For let $a < b_1 < \dots < b_n = b$ be a maximal chain in P , where $b^* \in \epsilon$. Assuming that the observation a remains valid for a length of time, there is a good chance that our observer will also make some of the observations b_1, \dots, b_m , thus updating ϵ' with some of the b_i . This will increase the chance of the excitation wave propagating all the way to b over *several* attempts at synchronizing memory with a recurring observation of a . Learning requires persistence.

We want to remark that all the above does not diminish the efficiency of the algorithm: though instead of an accurate result the de-idealized algorithm only produces an approximation, the improvement in the quality of the recorded data is a linear function of the running time. Finally, we would like to turn the attention of the reader to the fact that the described updating process is, in fact, both an updating and a retrieval process. It seems natural to identify these two aspects of database management for a living organism.

3.4. Other aspects of the proposed realization. We believe that the above analysis provides sufficient grounds for asking the question whether it is possible to realize a database structure as above using a neuronal network. If the answer is affirmative, then one should consider the way in which the other components of an observer \mathcal{O} could be realized by such a network. Here are some speculations in this direction:

Realizing the excitation function p : Here is a naïve and partial approach to encoding the excitation function p , motivated by the idea of propagation of excitations through p : a neuron for a can be viewed as supplying charge (through its axons) to every immediate successor $b \in P$. In order to force b to fire as a consequence of a firing, one needs to balance two parameters of this subsystem: (1) the action potential of the neuron b and (2) the amount of charge delivered to the dendrites of b through its connection with a . Tweaking this pair of parameters for every such pair (a, b) will affect the range over which any excitation wave can propagate. Such a realization would also provide a tool for interpolating between ‘rigid binary thinking’ and ‘probabilistic thinking’, and motivates replacing the boolean algebra \mathcal{B} (appearing as the range of the realization map f in the definition of an observer) by some other algebra, realizing other kinds of logic.

Structural updating: Neurons in brains of living organisms were observed in the process of creating/destroying synaptic connections while the observed animal was solving a problem. The ideas presented above poises the question about how restructuring of P can be achieved through tweaking the excitation parameters of the system as we had just described. The next section discusses the logical aspects of the problem formally, but the question remains whether neuronal networks can be used to realize the proposed database management model, and even more importantly – which cognitive phenomena can be simulated using such structures and how deep can one proceed with this analogy?

4. STRUCTURAL UPDATING OF OBSERVERS.

4.1. A deformation space for observers. There seem to be four components to update in an observer $\mathcal{O} = (P, f, p, \epsilon)$. We have already discussed possible ways for altering ϵ in the context of the basic search problem, but that discussion required the structural component of the database – the poc-set P – to remain constant. We must now address the problem of updating P .

Altering P may involve the formulation of new questions, or a new realization that certain questions should be made comparable (or even equal) after being considered as transverse over some initial time period. Such alterations to P will result in a redefinition of all the components of \mathcal{O} . In just a few words, our idea is that the trigger for restructuring P should be a significant change in the excitation function p : over time, some of the perceived events become negligible, and a clean-up operation is called for, to get the observer rid of the trouble of maintaining any record of such events. Changes in the values of the excitation function are external to the algorithmic structure of memory: these we assume to be generated by the physical organism, as a direct reaction to

observations made by the sensors it possesses.

Structural alterations should occur as a result of sensors being added, increased precision in existing sensors, communication with other observers and other observations causing a re-evaluation of excitation levels. In other words, alterations of the poc-set structure underlying the memory graph should correspond to a change in the natural language used by the observer. More precisely, we note that distinct observers may share common questions and be well aware of this fact. A formal way of expressing this in a general context is to have all questions ‘tagged’ by symbols from a prescribed alphabet \mathbb{A} , and all the participating agents using the same set of tags.

It is convenient for \mathbb{A} to be an infinite set. For every subset A of \mathbb{A} , let $D(A)$ denote the set of symbols $A \sqcup A^* \sqcup \{0, 0^*\}$ endowed with the complementation operator given by $a \mapsto a^*$ for all $a \in A \sqcup \{0\}$ and $a^* \mapsto a$ for all $a \in A^* \cup \{0^*\}$. Let $\mathcal{P}(A)$ denote the set of all poc-set structures of the form $(D(A), \leq, *)$ for which the symbol 0 is the minimum element. By $\mathcal{P}(\mathbb{A})$ we mean the union of $\mathcal{P}(A)$ over all finite subsets A of \mathbb{A} . For $P \in \mathcal{P}(A)$, we shall say that $D(A)$ is the *support* $\text{supp}(P)$ of P .

Denote by Δ^P the set of all functions $q : V\Gamma(P) \rightarrow [0, 1]$ satisfying $\sum_{u \in V\Gamma(P)} q(u) = 1$. Δ^P is a standard Euclidean $(|V\Gamma(P)| - 1)$ -dimensional simplex. To any pair (P, p) with $P \in \mathcal{P}(\mathbb{A})$ and $p : V\Gamma(P) \rightarrow [0, \infty)$ non-zero, we associate the point

$$\text{Pr}_{\mathcal{O}}(p) = \frac{p}{\sum_{u \in V\Gamma(P)} p(u)} \in \Delta^P.$$

Let us consider what happens as $q \in \Delta^P$ is being moved toward the boundary of Δ^P : in the eye of an observer \mathcal{O} with underlying poc-set P and $\text{Pr}_{\mathcal{O}}(p) = q$, the events corresponding to diminishing values of q gradually become negligible; in the limit (as p reaches a face F of Δ^P), \mathcal{O} will ignore such events, treating them as irrelevant. As a result, some questions may become redundant in the eyes of \mathcal{O} or the relations between them might change. However, we note that approaching the boundary of Δ^P from inside Δ^P will never result in a pair of nested elements of P becoming transverse, while a pair of transverse elements may transform into a nested pair or even into a pair of equal/complementary elements. The resulting ‘degenerate’ poc-set should then correspond to the face F . A convenient notion in this context is:

Definition 4.1 (Corners). Suppose P is a finite poc-set. A *corner* of $\Gamma(P)$ is the full subgraph of $\Gamma(P)$ induced by a set of vertices of the form $V(a, b)$ for some proper elements $a, b \in P$.

Putting the preceding discussion in other words, a degeneration of P should occur whenever $\text{Pr}_{\mathcal{O}}(p)$ assigns a negligible value to some corner of $\Gamma(P)$. Here is a very simple example to keep in mind:

Example 4.2. Recall example 2.19, where we had two poc-sets P and Q each generated by a pair of distinct proper elements a, b so that $a \pitchfork b$ in Q and $a < b$ in P . Then $\Gamma(Q)$ is a 4-cycle, while $\Gamma(P)$ is an interval with three vertices – see figure 8, left-hand side.




FIGURE 8. A simple example of degeneration of a memory graph due to recalculation of probabilities.

Assigning equal excitation probabilities to the four vertices of $\Gamma(Q)$ places the corresponding weighted graph at the barycenter of Δ^Q . Now consider a perturbation of those probabilities as described in the central diagram of figure 8: the excitation probability of the corner $V(a, b^*)$ vanishes as ϵ approaches zero, leading to a restructuring of the poc-set Q into the poc-set P .

Definition 4.3 (Degeneration). If $P, Q \in \mathcal{P}(\mathbb{A})$, say that Q degenerates into P (denoted by $P \leq Q$) iff $\text{supp}(P) \subseteq \text{supp}(Q)$ and there exists a *retraction* of Q onto P : a poc morphism $r : Q \rightarrow P$ restricting to the identity on $\text{supp}(P)$.

Using this notion, we are now able to glue all the simplices Δ^P , $P \in \mathcal{P}(\mathbb{A})$ to obtain a parameter space encoding the relationships among all the graphs $\Gamma(P)$ supported on \mathbb{A} :

Definition 4.4 (The Deformation Space of Observers over \mathbb{A}). For every $P, Q \in \mathcal{P}(\mathbb{A})$ satisfying $P \leq Q$ and for each retraction $r : Q \rightarrow P$, we define an injective map $\delta = \delta_r : \Delta^P \rightarrow \Delta^Q$ by setting

$$(7) \quad (\delta(p))(u) = \begin{cases} p(v) & \text{if } u = r^\circ(v) \\ 0 & \text{if } u \notin \text{Im}(r^\circ) \end{cases}$$

for every $u \in V\Gamma(Q)$. We then construct a space $\mathcal{U} = \mathcal{U}(\mathbb{A})$ — the *deformation space of observers over \mathbb{A}* — as the quotient of the union $\bigcup_{P \in \mathcal{P}(\mathbb{A})} \Delta^P$ by the equivalence relation generated by the relations of the form $p \equiv \delta_r(p)$ for all $p \in \Delta^P$, all $P, Q \in \mathcal{P}(\mathbb{A})$ and all retractions $r : Q \rightarrow P$.

For every $p \in \mathcal{U}$, denote by $\text{ch}(p)$ the unique simplex of \mathcal{U} containing p in its interior. Equivalently, $\text{ch}(p)$ is the top-dimensional simplex of \mathcal{U} containing p .

Thus \mathcal{U} is obtained from the collection of all simplices of the form Δ^P by identifying some of them along faces (gluing). The main characteristic of the space \mathcal{U} is that any path in this space corresponds to a process of altering some memory graph through introduction of questions (moving away from a face of simplex) or removal/identifications of/among redundant questions. Also note that any two simplices Δ^P and Δ^Q occur as

faces of Δ^R , where R is the poc-set with support $\text{supp}(P) \cup \text{supp}(Q)$ and no relations (every pair of proper questions is transverse). This shows that \mathcal{U} is the result of carrying out face identifications on one infinite-dimensional simplex.

4.2. The Updating Postulate. Up till now we have not made any assumption regarding the way in which the updating of an observer occurs, leaving the model static: at any point in time, an abstract observer can be associated with a real-life observer by, say, freezing the latter and studying the poc-set structure of the contents of its memory. Of course, it is the evolution of an observer that is of interest in relation to studying learning processes.

Consider the postulate: *If a sequence $\mathcal{O}_n = (P_n, f_n, p_n, \epsilon_n)$ ($n \geq 0$) represents consecutive stages in the evolution of the same observer, then, for every $n \geq 1$, either $\text{ch}(p_{n-1})$ is a face of $\text{ch}(p_n)$ or $\text{ch}(p_n)$ is a face of $\text{ch}(p_{n-1})$.*

This postulate means, essentially, that, whatever the implementation of any given observer, and whatever procedure is used for updating it, any such updating results either in a degeneration as described above, or in an ‘inverse degeneration’. Thus, one could imagine the evolution of a memory graph to be an alternating sequence of expansion/contraction moves.

The reason for stating such a postulate is that of economy: an attempt should be made to free the observer of the burden of maintaining low-priority vertices in $\Gamma(P_{n-1})$ whenever the same information can be encoded by degenerating P_{n-1} into a poc-set P_n with a smaller dual graph. Since $\Gamma(P)$ tends to grow exponentially with the size of P , this optimization becomes necessary due to its potential for conservation of resources.

The point of view described above provides a possible explanation of why humans, for example, find it so hard to update a set point of view (on practically any given issue), or, more generally, why good studying of a recurring phenomenon results in automatisms that often bar the student from creatively responding to an unpredicted change of circumstances. Looking again at the example in figure 8, it is easy to imagine the amount of destruction resulting from a more complex degeneration process (occurring in a more complex graph). It is then stunning to see how simple it is to achieve this kind of update (e.g., given an implementation such as the one proposed in section 3, this kind of update can be achieved through rewriting just a few pointers and then collecting the garbage), compared to how hard it is (algorithmically) to build a new hierarchy of related questions, perhaps reconstructing parts of the old one, and then shaping the graph (through continued observation leading to degenerations) so that it fits a more complete picture of the observed reality.

While the price to pay for mistakenly erasing low-priority nodes is great, the benefit of eliminating objective redundancies may outweigh this price if the threshold for erasing a low priority vertex is low enough.

To summarize, we have just described how the updating postulate turns the processes of updating memory, logical deduction and forgetting into different aspects (depending

on context) of the same principle, by restricting the effect of an ‘elementary updating process’ on the memory structure.

4.3. Structure of \mathcal{U} vs. Natural Language. Let us discuss an additional aspect of the idea of degeneration in the memory graph of an abstract observer. Suppose now that $\epsilon > 0$ is given with the property that every observer $\mathcal{O} = (P, \dots)$ in a given group of observers considers a vertex $u \in VT(P)$ negligible if $\Pr_{\mathcal{O}}(u) < \epsilon$. Once again, let \mathbb{A} be the alphabet of questions recognized by the members of the group.

For any $n \geq 0$, let $\mathcal{U}^{(n)}$ denote the n -skeleton of $\mathcal{U} = \mathcal{U}(\mathbb{A})$ – the union of all simplices of \mathcal{U} having dimension n or less.

Then, whenever $|VT(P)|$ exceeds $\frac{1}{\epsilon}$ (for some $P \in \mathcal{P}(\mathbb{A})$), a whole ball about the barycenter of the simplex Δ^P becomes irrelevant for the discussion of this particular group of observers. This makes the $\frac{1}{\epsilon}$ -skeleton of \mathcal{U} much more relevant to the discussion of ‘language’ in the given group than the space \mathcal{U} itself. Moreover, replacing \mathcal{U} by $\mathcal{U}^{(1/\epsilon)}$ in the role of a ‘space of all observers’ increases its topological complexity, which may also be relevant in the discussion of language structure and formation.

Here is a more precise formulation of what we mean by this. Consider the process of parents teaching their newborn child. The newborn has many ‘questions’ available, but not much can be made out of them at the very beginning: they are not ordered yet. The environment provides a lot of input, and the role of the parents is to serve as a filter, protecting the child from input it is yet incapable of processing. The result is that the parent slowly synchronizes the child’s memory structure with their own, increasing the chances that the child’s sensors carry the same meaning as analogous sensors in the parent. As a result, the individuals in a stably evolved population will have many common patterns in their memory structures: a vast majority of individuals will agree on certain associations between different inputs (consider our attaching a very specific meaning to the sound of the word ‘green’, unless we are color-blind). Thus, for any given population, a subset \mathbb{B} of \mathbb{A} exists, for which a poc-set structure is already decided among the adults of that population, and all discourse and synchronization among adults is confined to the sub-complex of $\mathcal{U}(\mathbb{A})$ – denote it by $\mathcal{U}(\mathbb{A}, \mathbb{B})$ – constructed using only those $P \in \mathcal{P}(\mathbb{A})$ for which the poc-set structure on $\text{supp}(P) \cap \mathbb{B}$ inherited from P coincides with the poc-set structure inherited from \mathbb{B} . The evolution of the population’s language is then partially described by the evolution of $\mathcal{U}(\mathbb{A}, \mathbb{B})$ over time; the learning goal of a child in the population is to synchronize their memory structure with that defined by the parent population.

In subsection 2.4 we discussed the effect of the capability for employing logical (and perhaps other) connectives for enriching the poc-set structure of an observer by creating ‘new questions from old’. The main observation was that an observer needs to balance between two extremes: over-using this capability leads to exponential inflation of the memory structure making it impossible to maintain efficiently; not using it enough prevents the memory structure from attaining a sufficient level of refinement and blocks the observer from capturing more subtle information about the observed system.

It is easy to see how this idea gets incorporated into the discussion of the deformation space $\mathcal{U}(\mathbb{A})$: simply impose additional algebraic structure on $\{0, 0^*\} \cup \mathbb{A} \cup \mathbb{A}^*$. For example, here is a natural way of adding Boolean connectives to the alphabet. For each $a, b \in \mathbb{A}$ one assumes that the symbols of the form $a \wedge b$, $a \vee b$ (and all incident finite formulae) are in the alphabet, together with all the ensuing relations, and we require all admissible poc-set structures to be synchronized with these symbols – e.g., for all $a, b, c \in P \in \mathcal{P}(\mathbb{A})$ require that $a \wedge b < a$, $a \wedge b < b$ and $(c < a) \wedge (c < b)$ must imply $c < a \wedge b$. Yet again, the result is a sub-complex of the original $\mathcal{U}(\mathbb{A})$. The two refined constructions we had just discussed can be combined together or compared – any option will yield an aspect of the phenomenon one can only call by the name *the language spoken by the population*.

5. DISCUSSION

We have defined and, to an extent, analyzed a family of databases designed to maintain the memory of an arbitrary entity observing the evolution of an environment by means of binary sensors. Motivated by the example of living organisms, we assumed the following:

- The observer evolves with time, possibly ‘growing’ new sensors;
- The observer may interact with the environment;
- The observer has mechanisms to assign excitation levels to observed data – we refer to these collectively as the observer’s *evaluation mechanism*.

Our main goal was to construct a memory system capable of dealing with arbitrary input, given the guidance provided by the evaluation mechanism. This means input is not treated according to its objective significance to the observer, but rather according to the *subjective* importance assigned to the input by the evaluation mechanism. When the evaluation mechanism and reality are in tune with each other, input from the environment contributes to a more accurate record of events in the database, and to better decision-making on the part of the observer.

5.1. The Invariance Principle and alternate logics. A major trait of our model is a clear separation of the algorithmic side (the database) from the physical side (the realization). On the former side, an observer $\mathcal{O} = (P, f, p, \epsilon)$ is required to maintain the poc-set structure P , the (conjectural) current state ϵ and the graph $\Gamma(P)$ with vertices weighted by p . On the latter side, we have the poc-morphism $f : P \rightarrow \mathcal{B}$ (recall \mathcal{B} is the Borel σ -algebra on X) *objectively* representing the observer’s sensors. Together, these components produce a picture of the mind of a real-life observer, frozen in a given moment of its evolution. The evaluation mechanism and actual physical realization of this system are given the task of effecting the transition of \mathcal{O} from any given current state to its next state. This is done through:

- (1) Sampling sensory input. In this context, the physical realization of the observer is considered a part of the observed environment.
- (2) Updating the current state record given by ϵ to fit the most recent observations (see section 3)
- (3) Re-evaluating excitation levels relevant to the current observations.

- (4) Optimizing the database structure $\Gamma(P)$ to fit the new excitation levels.

Note how clear cut is the division of labor between the algorithmic and physical components of the model: stages (1) and (3) are totally dependent of the physical organism, while stages (2) and (4) may be formulated completely in terms of the algorithmic structure, only depending on the physical ‘shell’ for their input and execution of the algorithmic tasks at hand.

In section 1.1 we had already mentioned the more obvious potential of such a ‘two-lobe’ approach. In a manner of speaking, this approach achieves a separation of ‘software’ from ‘hardware’, enabling the study of the question: what kind of hardware is capable of supporting this software?

Note that the representation map f is the only objective component of an observer \mathcal{O} . The choice of range for f (the Boolean algebra \mathcal{B}) corresponds to the assumption that \mathcal{O} uses binary sensors to observe its environment. As a result, a deeper aspect of the invariance principle comes to light: one could weaken the binarity assumption considerably by replacing the range of f by a larger algebra quantifying over X . For example, one could use the space $\Psi(X, \mu)$ of probability distributions on (X, μ) as a range for f . Indeed, the space $\Psi(X, \mu)$ has a fuzzy logic underlying it, while it still retains a poc-set structure induced from \mathcal{B} : for every $\psi, \psi' \in \Psi(X, \mu)$ one defines $\psi^* = 1 - \psi$, and $\psi \leq \psi'$ iff $\psi(x) \leq \psi'(x)$ holds throughout X . This change opens new horizons we have not explored yet, as it greatly alters the interpretation of the database. To be precise: the observer \mathcal{O} *observes* fuzzy phenomena, but *records* them as if they were deterministic.

Quantum logics and temporal logics are of particular interest to us in this context: replacing \mathcal{B} by a quantum logic may be expected, according to [Kau], to enable \mathcal{O} to consider self-referential statements; a logic with a temporal structure could allow \mathcal{O} to become naturally aware of the passage of time, and consequently – aware of its own learning processes.

At the same time, while the range of the representation map is allowed to vary, it is not necessary to change the structure of the database itself. It is then the *interpretation* of the database structure and its role in the interaction between the database and the physical realization that changes when the binary logic of our model is replaced by another logic. This opens the door to questions regarding the nature of our own (e.g., human) ‘guessing algorithms’, and motivates one to ask whether a memory structure such as the one we are proposing in this paper can be used to enforce decision making by an autonomous agent in non-deterministic situations.

5.2. Explanatory power of the model. The introduction strongly emphasizes the idea that frequently observed weaknesses of the human thought process, including the destructive process of forgetting, should serve as testing stones for any model contending for the high title of “model of the memory of living organisms”. More precisely: any contender must be capable of demonstrating as many such phenomena as possible to be different aspects of its *normal* operation.

Now that all the proper language has been developed and all modeling assumptions have finally been stated, we can summarize the most notable achievements of our model

on this front. Put in our new language, the memory structure of an organism in an \mathbb{A} -speaking population is modeled by a sequence $(\mathcal{O}_n)_{n \in \mathbb{Z}}$ of observers $\mathcal{O}_n = (P_n, f_n, p_n, \epsilon_n)$. Recall that structural updating is carried out in accordance with the updating postulate from 4.2, so for every $n \geq 1$ we have either that $\text{ch}(p_{n-1})$ is a face of $\text{ch}(p_n)$ or that $\text{ch}(p_n)$ is a face of $\text{ch}(p_{n-1})$ in the deformation space $\mathcal{U}(\mathbb{A})$.

Permanent loss of information. There are two ways in which information is lost by our database structure. Permanent loss of information occurs whenever the transition $\mathcal{O}_n \rightarrow \mathcal{O}_{n+1}$ involves a degeneration of a corner $V(a, b)$ of $\Gamma(P_n)$ whose representation in \mathcal{B} has non-zero probability: our observer will, never the less, consider the logical implication $a < b^*$ as a property of X unless renewed observations somehow force a reverse structural update. There are two important observations to make in this context:

- (1) ‘Forgetting’ the event in \mathcal{B} corresponding to the corner $V(a, b)$ is equivalent to committing the statement $a < b^*$ to memory. Therefore, every act of memorization – every act of recording a conclusion about the structure of X , if you like – may, potentially, result in loss of information. However, structural updating is necessary for keeping the database from blowing up, so this price has to be paid whether we like it or not.
- (2) Information lost (by the database) in this way is very hard to recover if we realize the database in a manner such as the one described in section 3. The reason is that – at least in the ideal implementation – recording the relation $a < b^*$ implies that whenever the observation a is made, propagation of excitation implies the observation b^* is made as well. In the non-ideal implementation, it is possible that the observation b^* will not be made if there are sufficiently many $x \in P$ satisfying $a < x < b^*$: dissipation of the excitation signal may cause the search algorithm not to reach b^* at all. In either case, an incoming observation of b and a may be simply discarded by the system unless the evaluation mechanism gives this input an excitation value that is too high to ignore. Possibly, this issue may be better understood in the context of a more realistic implementation of the database by, say, a neuronal network.

Temporary loss of information. We consider situations when the information sought for is contained in the database, but, for some reason, is hard to find. Formally speaking, the system is being fed an observation $a \in P$ (or observations $a_1, \dots, a_n \in P$), and is expected to produce the conclusion $x \in P$ – to respond to the observation in a way that provides proof (to the environment, in some sense) that it has ‘understood’ the implication ‘ a implies x ’; however, the desired demonstration fails to occur.

In the context of an implementation based on propagation of excitation, these situations are easy to explain/recreate: an input, or combination of inputs, triggers a cascade of ‘internal’ observations, due to propagation of the excitation signal(s); if the signal dissipates before the node x is reached, then the observer will fail to produce the required reaction. However, repeated queries along a path (in P) from a to x may result in the excitation wave reaching x *eventually*.

Bound on parallel processing. This is a point we have hardly touched in this paper, but would like to point it out as a motivation for future research. If the network of nodes

realizing P can be excited at several nodes a_1, \dots, a_n simultaneously, corresponding to the observation $a_1 \wedge \dots \wedge a_n$ being made, then the range of reachable conclusions potentially increases with n . The observation $a_1 \wedge \dots \wedge a_n$ serves as proof that $\{a_1, \dots, a_n\}$ is a coherent family, which means $V(a_1, \dots, a_n) \subset \Gamma(P)$ is non-empty. If there is some *a priori* upper bound on n (an upper bound on the number of simultaneous observations), then, as the set of available sensors grows, the ability of the observer to relate to specific vertices of $\Gamma(P)$ diminishes unless the organism has a tool for introducing sensors corresponding to conjunctions of families of other sensors. In any case, a bound on the ‘bandwidth’ of the implementation of the poc-set P automatically implies a restriction on the ability of the organism to deal with complex input.

Some organisms, like modern humans, have developed tools to help them circumvent such difficulties. We are able to articulate the results of our thought process and write them down, thus creating an artificial feedback loop with its own memory capacity. Reading our text at a later time re-excites thought processes that were abandoned earlier, which, armed with new data, updated memory and new notions may now have a better chance of reaching the desirable goal.

Noise. Remember that song that you cannot stop humming? The gossip information from television that you recall with an ease igniting the envy of all the material you have read for work and successfully forgotten? How about the dead silence, or the chirping birds, or deafening heavy metal music that you need for concentration?

Our model explains these through ‘bandwidth’ considerations from the preceding paragraph and through known properties of the evaluation mechanism in living organisms. For us and for many other animals, the most significant component of our excitation levels seems to be physiological stress.

Physiological stress plays the same role in biology as temperature in physics: when stress is totally absent, the animal is overly docile and disinterested with its environment; when there is too much stress, metabolic reactions peak and the animal becomes incapable of reacting to the environment due to the total chaos in the input it has to process. Depending on the processing mechanism, some inputs (normally weak, periodic, predictable) will stress an animal just enough to cause it to ignore most common but irregular minuscule sources of stress (distractions), allowing it to use the rest of its capacity for parallel processing for quality ‘thinking’ (this is a state of ‘calm’ for the animal). Other inputs get evaluated so high, that they keep occupying the animal’s computational resources for long periods of time – at least as long as the inputs persist – which results in wasting computational resources on the evaluation and re-evaluation of those signals and any ‘memories’ they may trigger as a result of the searching process (as described in section 3). Reliving these memories, in turn, may result in more stress, creating a cycle that only breaks with the introduction of a much more powerful signal or with the disappearance of the majority of the stressing factors from the observer’s horizon.

This is, of course, a speculative picture, but we find the simplicity of the explanation appealing. In particular, our model then serves as a (formal, mathematical) link between the algorithmic and psychological aspects of learning.

5.3. Biological connections and natural language. The preceding review started with computational aspects of information processing (as realized in the proposed model) and ended with the biological ones. In the context of living beings (which are forced to interact with their environment), neither can be discussed separately from the notion of a natural language. Different organisms have evolved their information processing tools to different levels (different parts of the human nervous system may be traced back to different stages of the planetary evolutionary process), and the same can be said about the evolution of their means for communicating with the environment. Thus, inevitably, the evolution of natural languages is related to the evolution of physical realizations of memory (e.g., brains) and of the algorithmic tools maintaining the underlying database structures.

The current textbook definition of a natural language is somewhat lacking in mathematical rigor. Definitions such as “by a natural language we mean human languages such as English, Spanish, Arabic etc.” are generally accepted, but can hardly be considered rigorous.

We believe our approach has a new bearing on the formal understanding of the notion of natural language, as well as the process of language acquisition at the level of the individual learner. Niyogi (see [Niy06] section 1.1, p.16) provides a learnability argument in favor of restricting the family of languages human learners are capable of producing:

“The necessary and sufficient conditions for successive generalization by a learning algorithm has been the topic of intense investigation by the theoretical communities in computer science, mathematics, statistics and philosophy. They point to the inherent difficulty of inferring an unknown target from finite resources, and in all such investigations, one concludes that *tabula rasa* learning is not possible. Thus children do not entertain every possible hypothesis that is consistent with the data they receive but only a limited class of hypotheses. This class of grammatical hypotheses is the class of possible grammars children can conceive and therefore constraints the range of possible languages that humans can invent and speak.”

The mathematical nature of these constraints remains unclear unless a reasonable formal universal model of a learner is provided. In some simplified sense, this is precisely what our model is. Moreover, we show that it comes equipped with a tool for formally defining the learning goal for language acquisition. In what follows, we will try to substantiate this claim using the ‘deformation spaces of observers’ we have recently considered in section 4.

In the context of humans, the notion of a natural language is directly associated with the use of sound for the articulation of one’s thought process. However, once the use of signs and signals is permitted – written language, for example – one has to face a growing number of alternatives stretching the definition far beyond the realm of the audible. For example, where should one place sign language? by what means should one communicate with the little green deaf visitors from Mars? If one is to take into account

so many variations on the original notion of natural language, then surely this notion should apply to every self-sufficient coherent system of signals allowing the exchange of information among sentient beings.

Stop. At this point, it will do us good to realize our inability to discern sentient beings from non-sentient ones. At best, we are able to distinguish solitary species from social ones. But perhaps this is where one finds the key to solving our problem: while *any* repetitive pattern of information exchange among physical entities may be considered a language (in a broader understanding of the term), we should be seeking a notion that is inherently linked with shared territory and ordered communication among organisms. It is no big surprise that humans and human languages have evolved from social primates rather than from solitary wasps, say. After all, from the evolutionary point of view, populations of solitary species have no use for complex articulated information exchange, while populations of social organisms may benefit from this ability.

It is this last idea that motivated the construction of a deformation space $\mathcal{U}(\mathbb{A})$ of \mathbb{A} -speaking observers. The space $\mathcal{U}(\mathbb{A})$ is best understood in the context of a population \mathcal{P} of organisms sharing a set \mathbb{A} of statements about X . The elements of \mathbb{A} should not be thought of as corresponding to words in the language spoken by the population, but rather to the *shared meanings of complete sentences*. An outside observer trying to study \mathcal{P} may initially be oblivious of the different meanings of the elements of \mathbb{A} , but they could try to guess those meanings from contextual data collected while observing the evolution of memory structures of members of \mathcal{P} . Similarly, a new member of \mathcal{P} – a newborn child, say – has to observe other members of \mathcal{P} and communicate with them in order to uncover the actual meaning of each element of \mathbb{A} .

Now, as we have already pointed out in section 4.3, shared meaning implies shared logical structure, so it is reasonable to assume that the adult part of the population will have a shared record of certain relations among elements of $D(A) = \{0, 0^*\} \sqcup \mathbb{A} \sqcup \mathbb{A}^*$, which, inevitably will show up in their memory structures. The young of \mathcal{P} can then be considered as trying to restructure their memory graphs accordingly. It is also plausible to expect various connectives to appear as tools for enriching the set of statements meaningful to the population: this introduces an algebraic structure on $D(A)$, e.g. conjunction and implication operators in the case of human languages. More generally, *any* shared structuring of the information exchange among members of \mathcal{P} ends up as a shared property of the adult memory structures that has to be learned (through restructuring) by the younger members. We have shown in 4.3 how additional structural properties of $D(\mathbb{A})$ (or subsets of the form $D(\mathbb{B})$, $\mathbb{B} \subset \mathbb{A}$) define special subspaces $\mathcal{U}(\mathbb{A}, \mathbb{B})$ of the deformation space $\mathcal{U}(\mathbb{A})$ with non-trivial topology derived from the properties of \mathbb{B} . Mathematically speaking, studying the topology and geometry of these subspaces corresponds to studying the language ‘spoken’ by \mathcal{P} , from an outsider’s viewpoint. From an insider’s viewpoint, a learner’s goal is to have its memory structure synchronized with that of the entire population: the population provides input consistent with \mathbb{B} while the learner’s evaluation mechanism provides guidance directing the evolution of the learner’s memory structure towards $\mathcal{U}(\mathbb{A}, \mathbb{B})$; inadequate guidance (e.g., the learner was born autistic) may prove insufficient for the learner to achieve the desired result;

inadequate pressure by the parent population (input contradicting the learner's evaluation mechanism) may result in failure to acquire the language of the population.

At the current stage, this last observation is not of much use to a linguist studying a specific language. However, the novelty in our point of view is in that we have shown how language acquisition occurs (in the framework of our model) as a result of a learning process that is *not* guided by a specific pre-defined problem, but by the very general problem of the learner (e.g., a child) attempting coherent communication with its immediate environment. The only guidance present is the *subjective* guidance by the evaluation mechanism, and this can be swayed in many different directions by an attentive teacher.

To summarize, all the above suggests that the constraints Niyogi discussed may be derived from two sources. The first is an accurate description of the evaluation mechanism providing the guidance for the learning process. The second is a good description of the logic employed by the learner and the way in which it is realized by the updating algorithms of the learner's memory structure. While our model assumed classical Boolean logic, our discussion of the invariance principle and its realization in our model shows how this assumption can be lifted. We leave this direction to future research.

REFERENCES

- [Gur] Peter J. Guralnik, *The role of aggression in establishing and maintenance of group structure in animal populations: Socium vs. compound*, Preprint, <http://www.math.ou.edu/~dguralnik>.
- [HP98] F. Haglund and F. Paulin, *Simplicité de groupes d'automorphismes des espaces à courbure négative*, The Epstein birthday schrift, Geometry and Topology Monographs, vol. 1, 1998, pp. 181–248.
- [Kau] Louis Kauffman, *Laws of form - an exploration in mathematics and foundations*, Book in progress, <http://www.math.uic.edu/~kauffman/Laws.pdf>.
- [Nic04] Bogdan Nica, *Cubulating Spaces with Walls*, Algebraic & Geometric Topology **4** (2004), pp. 297–309.
- [Niy06] Partha Niyogi, *The computational nature of language learning and evolution*, Current Studies in Linguistics, MIT Press, 2006.
- [Rol98] M.A. Roller, *Poc sets, median algebras and group actions*, University of Southampton, Faculty of Math. stud., preprint series (1998).
- [Sag95] Michah Sageev, *Ends of groups pairs and non-positively curved cube complexes*, Proc. London Math. Soc. **3** (1995), no. 71, 586–617.
- [Sha48] C.E. Shannon, *A mathematical theory of communication*, The Bell System Technical Journal **27** (1948), 379–423, 623–656.

DEPT. OF MATHEMATICS, UNIVERSITY OF OKLAHOMA, NORMAN, OK 73019
E-mail address: dan.guralnik@ou.edu